

AMENDMENTS TO THE SPECIFICATION

Please replace paragraph [0010] with the following amended paragraph:

[0010] Resilient packet ring system 100 includes a number of ring stations (station 0 130, station 1 140, station 2 150, . . . and station N 160) interconnected by a ring structure utilizing unidirectional, counter-rotating ringlets. Each ringlet is made up of links between stations with data flow in the same direction. The ringlets are identified as ringlet0 110 and ringlet1 120. This standard allows a data frame to be transmitted on either of the two connected ringlets. For example, a unicast frame is inserted by a source station and copied by the destination station. For efficiency, the destination also strips the now irrelevant stale frame. The portion of a ring bounded by adjacent stations is called a span, and thus a span is composed of unidirectional links transmitting in opposite directions. The RPR dual-ring topology ensures that an alternate path between source station and destination station(s) is available following the failure of a single span or station. Fault response methods include pass-through and protection[[.]], as described in the standard.

Please replace paragraph [0029] with the following amended paragraph:

[0029] The systems, methods, apparatus and software of the present invention can be implemented in the context of network transport devices designed to be compliant with the IEEE 802.17 Resilient Packet Ring (RPR) standard, which is ~~underdevelopment~~ under development as of the filing of the present application. The current version of the standard is described in IEEE Draft P802.17/D2.2, *Resilient Packet Ring (RPR) Access Method & Physical Layer Specifications*, April 9, 2003, which is hereby incorporated by reference herein in its entirety. However, the systems, methods, apparatus and software of the present invention need not be limited to RPR implementations. In general, the systems, methods, apparatus and software of the present invention can be utilized in the context of a variety of different networking structures and topologies.

Please replace paragraph [0041] with the following amended paragraph:

[0041] Although not illustrated, there is typically a fairness instance associated with each of the two ringlet datapaths. A fairness instance can be uniquely identified by the combination of the identity of the station, usually determined by the MAC address of the **stations station's** MAC device, and the identity of the ringlet carrying the data traffic whose rates are regulated by the fairness instance. In some cases, e.g., when a station is in the center-wrapped state as described by the RPR standard, the station includes a single fairness instance associated with both ringlets. A fairness instance computes a locally significant fair rate value that provides the basis for the computation of other rates that are communicated to other stations on the ringlet, to the MAC datapath, and to the MAC client. An administrative weight is assigned to each fairness instance to permit the scaling of fair rate values among stations on the ringlet. This allows one station to use a larger share of available capacity than another station without violating fairness principles. The capability of allowing such scaling is known as weighted fairness. A rate communicated from one station to another is typically normalized in order (1) to ensure that the rate is uniformly interpreted by stations on the ringlet and (2) to scale the rate value to allow it to be efficiently encoded as an integer value within a fair rate field of a fairness frame. Additionally, ramping describes the gradual increase or decrease of a rate. Fairness procedures typically make use of several ramping algorithms. Ramping methods usually employ a ramping coefficient that can be configured to increase or decrease the degree of ramping. The various algorithms and techniques used in support of RPR fairness are described in greater detail in the RPR standard. The systems, methods, devices, and techniques described herein can make use of and generally be used in conjunction with these fairness techniques. Moreover, the systems, methods, devices, and techniques described herein can further implement and/or utilize fairness algorithms and techniques not described in the RPR standard.

Please replace paragraph [0053] with the following amended paragraph:

[0053] Packet processor 520 typically operates in conjunction with some manner of additional buffering such as transmit/receive buffer 530. Transmit/receive buffer 530 facilitates packet processing by providing temporary data storage to packet processor 520. Such buffering can allow packet processor to more efficiently perform its primary task of routing data to an appropriate network device. For example, transmit/receive buffer 530 facilitates the implementation of certain congestion avoidance algorithms used to manage the data queue depth/latency. Examples of such algorithms include the random early detection (RED, sometimes also referred to as “Random Early Drop” or “Random Early Discard”) algorithm and the weighted RED algorithm which ultimately determine if a packet should be ~~enquired~~ enqueued as requested, or dropped. Other algorithms and techniques can make use of buffering and queuing devices such as transmit/receive buffer 530.

Please replace paragraph [0058] with the following amended paragraph:

[0058] In addition to VDQ shapers 570, link shapers 580 can also be implemented. Whereas VDQ shapers 570 shape transmission rates based on advertised MAC client receive rates (or some other similar information), link shapers 580 adjust data transmission based on bandwidth limitation of particular network links. In the example of Figure 5, there are three link shapers illustrated, one corresponding to each of the three links of a single ringlet of network 400 upon which data from station 0 would be transmitted. For example, it is assumed that data transmitted by station 0 would be destined for one or more of station 1 420, station 2 430, and station 3 430. Since station 0 will typically not transmit data over the ring network to itself (although in some embodiments that can be supported with one or more additional VDQs, VDQ shapers, and/or link shapers), data transmitted by station 0 on a particular ringlet will traverse those links that do not end at station 0. Therefore, using ringlet 0 401 as an example, data transmitted by station 0 410 will only traverse links associated with spans 415, 425, and 435. Since data destined for station 1 420, i.e., data dequeued from VDQ_S2 563, need

only traverse one link, it passes through ~~as~~ **[a]** single link shaper associated with that link. Data destined for station 2 **430** traverses both the previously described link and the corresponding ringlet's link in span **425**, and so it passed through two link shapers, and so on.

Please replace paragraph **[0060]** with the following amended paragraph:

[0060] As illustrated, link shapers **580** correspond to the links associated with a single ringlet, and thus there may be the need for ~~additionally~~ **additional** link shaper instances corresponding to the links of the other ringlet. In another embodiment, span shapers are used instead of link shapers. Span shapers can take the bandwidth restriction of complete ring spans into consideration, thereby eliminating the need for separate shapers for each of the two MAC datapath instances.

Please replace paragraph **[0064]** with the following amended paragraph:

[0064] Additionally, a variety of different techniques can be used to determine the fair rates, ramp factors, or other information indicating the need to reduce traffic targeting a particular MAC client. These techniques may have equal applicability to any of the systems, methods, devices, and techniques ~~deseried~~ **described** herein. In one example, described below, a conservative fairness algorithm determines new fair rates based on a number of factors including the extent to which a burst buffer or VDQ is full, receive rates for both fairness enabled and non-fairness enabled traffic, current fairness rates, and ramp coefficients.

Please replace paragraph [0067] with the following amended paragraph:

[0067] Similarly, as the buffer/queue capacity rises above the first threshold, i.e., from the hungry state into the satisfied state, ~~an~~ [[a]] new ramp factor can be calculated so as to lead to a decrease in traffic:

$$RampFactor = -\frac{RcvdFERate}{(RcvdFERate + FwdRate) \times RampCoeff},$$

where *FwdRate* is the rate at which fairness enabled traffic is transitted (forwarded) through the MAC device. Using this *RampFactor*, a new fairness rate is calculated as shown above. In situations that might be considered more critical, e.g., as the buffer/queue capacity rises above the second threshold, i.e., from the satisfied state into the full state, the amount of the ramp can be further increased multiplying the *RampFactor* by a constant greater than one, e.g., two.